

THE REINFORCEMENT OF LEAST-FREQUENT INTERRESPONSE TIMES¹

DONALD S. BLOUGH

BROWN UNIVERSITY

A new schedule of reinforcement was used to maintain key-pecking by pigeons. The schedule reinforced only pecks terminating interresponse times which occurred least often relative to the exponential distribution of interresponse times to be expected from an ideal random generator. Two schedule parameters were varied: (1) the rate constant of the controlling exponential distribution and (2) the probability that a response would be reinforced, given that it met the interresponse-time contingency. Response rate changed quickly and markedly with changes in the rate constant; it changed only slightly with a fourfold change in the reinforcement probability. The schedule produced stable rates and high intra- and inter-subject reliability, yet interresponse time distributions were approximately exponential. Such local interresponse time variability in the context of good overall control suggests that the schedule may be used to generate stable, predictable, yet sensitive baseline rates. Implications for the measurement of rate are discussed.

The rate with which an animal emits a simple response can clearly show the effect of stimulus and other variables upon behavior. But as a quantitative dependent variable, rate is often disappointing. Rates under a given schedule of positive reinforcement tend to vary from subject to subject, and from time to time. It is hard to reproduce a given rate in a given subject, upon returning to a given set of experimental conditions (*e.g.*, Ferster and Skinner, 1957, p. 369, 370; Sidman, 1960, p. 253). Sometimes rate seems to become insensitive to usually effective variables (Sidman, 1960, pp. 176-ff).

Such difficulties suggest that standard schedules do not control important factors affecting response emission. One such factor is interresponse dependence. When behavior is maintained by intermittent reinforcement, certain interresponse times tend to become highly probable, others infrequent. Patterns of successive interresponse times may emerge. Stereotyped unrecorded behaviors probably play a major role in mediating such dependencies.

Interval schedules favor the growth of interresponse dependencies by failing to control the behavior that comes just before a reinforced response, so that subjects tend to settle into superstitious patterns of behavior (Skinner, 1948) that yield unpredictable rates and patterns of response. Ratio schedules chain response tightly, yielding short, uniform interresponse times and insensitive rates. The differential reinforcement of low rate specifically favors interresponse dependence.

To the extent that responses are chained to preceding responses or to unrecorded prior behavior, stimulus or other variables may lose control over response emission. The ideal behavioral baseline would appear to be one that lacks such stereotypy, yet is statistically stable and reproducible. The schedule of reinforcement described here attempts to approximate such a baseline.

Physical systems that emit responses randomly in time have the desired characteristics. Radioactive decay is a classic example: particles appear at random intervals, yet the overall rate of emission is statistically stable. An ideal subject that behaved in this way would generate an interresponse-time distribution described by the exponential density function (Equation 1):

$$f(t) = \lambda e^{-\lambda t} \quad (1)$$

where λ represents mean rate and t represents

¹This research was supported in part by USPHS grants MH-02456 and MH-08355. Mrs. Patricia Blough, Miss Vicky Gray, Mr. Lloyd Marlowe, and Mr. Charles Shimp contributed to the research by subject running, computer programming, comment and criticism. Reprints may be obtained from the author, Walter S. Hunter Laboratory of Psychology, Brown University, Providence, Rhode Island 02912.

the time between responses (Mueller, 1950; McGill, 1963). The form of this function is shown by the idealized raw data of Fig. 1 ("x" symbols). The form is simplified by the logarithmic transformation, which yields a straight line of negative slope equal to λ , and the interresponse-times per opportunities transformation, which yields the constant λ (Fig. 1). These transformations are discussed and used below.

The present schedule operates to force the subject's interresponse-time distribution to approximate the ideal form just described. It is called "the reinforcement of least-frequent interresponse times" or the "LF schedule". Briefly stated, it favors local interresponse time variability by reinforcing unlikely interresponse times. To regulate responding in this way, responses must be continuously monitored, and reinforcements designed to adjust the interresponse-time distribution must be fed back to the subject at appropriate times. This was done by a LINC computer (Clark and Molnar, 1964), provided through an evaluation program sponsored by the National Institutes of Health.¹ This digital machine, running on line, programmed reinforcements, stored response data, and analyzed and displayed results. Most of the present graphs are photographs of LINC oscilloscope displays. Another application of the LINC to operant programming has already been described (Weiss and Laties, 1965).

METHOD

Subjects

Three White Carneaux pigeons were maintained at approximately 75% of free-feeding weight by supplementary feeding, if necessary, after each experimental session. Two birds had served in discrimination experiments involving variable interval schedules for many hours before this experiment. The third, 007, had 2 hr of continuous reinforcement and variable interval training before this study commenced.

Apparatus

The birds worked simultaneously in three standard Lehigh Valley pigeon chambers, which were dark except for a white-lighted response key. The switch on each key closed on application of about 14 g static pressure.

A loudspeaker in the chamber supplied white masking noise.

The LINC computer sensed closures of the key switches. Acting on rules described below, it programmed reinforcements and delivered them via banks of output relays. The LINC stored response data during a session and wrote it on magnetic tape at the end of a session. Subsequently, other stored programs recalled the data, analyzed it, printed out the results, and displayed them in graphic form on an oscilloscope.

Procedure

The experiment ran daily (with three or four exceptions) over a period of nine months. Each daily session lasted 80 min. Pecks on the response key produced occasional reinforcement consisting of 3.3 sec access to mixed grain. During reinforcement, the key light went out and a light over the food magazine came on. During a session, a bird might receive from 10 to 100 reinforcements, depending on program parameters in the following schedule.

The schedule was designed to reinforce responses that terminated a given subject's least frequent interresponse time. To determine interresponse time (IRT) frequency, it is necessary to collect IRTs occurring over some finite interval, or "IRT bin". The size of this IRT bin, and the point at which it falls on the IRT continuum, will of course be crucial in determining how many IRTs will fall into the bin. For most purposes—data analysis, for example—these bins are of equal size. Each bin might be a half-second long, so that responses falling in the first half-second after a response would go into the first bin, those in the second half-second in the second bin, and so on.

The crucial element in the schedule is the use of bins of varying size. As mentioned above, a process emitting responses randomly will generate an exponential distribution of IRTs (equation 1, Fig. 1 and 2). With equal bin sizes, the least frequent IRT would be the longest. However, the bin sizes may be adjusted in such a way that the ideal emitter, responding at a constant mean rate, will drop equal numbers of IRTs into each bin. This adjustment can then represent the fact that, for the ideal emitter, responses are equally probable at any moment—that their conditional probability, estimated by frequency per

opportunity (Anger, 1956) is constant. The bins are therefore chosen so that for a given mean rate λ they divide the area under the exponential IRT curve into equal parts. Figure 2 illustrates this subdivision. Here, the curve for $\lambda = 1/2$ has been marked off into 16 bins of equal area. The bin boundaries change with λ ; for $\lambda = 1$ the curve falls more steeply and the bins squeeze to the left, while for $\lambda = 1/4$ the bins expand to the right.

Figure 2 also shows that the ideal curve used to determine bin sizes does not start at IRT = 0, but at IRT = 0.8 sec. This was necessary because real subjects have a minimum IRT, and very short IRTs appear to have special characteristics, to be discussed later, that justify their exclusion here. The ideal IRT distribution thus becomes:

$$f(t) = \lambda e^{-\lambda(t-0.8)} \quad (2)$$

The IRT bins, determined for a given λ as illustrated in Fig. 2, controlled reinforcement. They will be called "reinforcement bins" to distinguish them from the conventional equal-sized bins used in data analysis.




Fig. 1. The "raw data" function ("x") represents the IRT distribution generated by an "ideal" subject emitting responses independently at a constant mean rate; the curve is exponential. The logarithmic transformation of these data is a straight line, with slope that varies with mean rate. The IRTs/OP transformation, defined in the text, indicates the conditional probability that a response will occur in any time bin; for exponential data, this probability is constant. The curves are arbitrarily displaced along the ordinate.

To determine which response should be reinforced, the LINC classified each incoming IRT into one of these reinforcement bins (or discarded it, if it fell below 0.8 sec). The machine kept a record of the 150 most recent classified IRTs from each bird², and also retained a running frequency distribution of IRTs by reinforcement bin. The reinforcement contingency for a given bird was reset at the time of reinforcement, on the basis of the IRT distribution at that moment. At this time, the reinforcement bins were scanned and the bin holding the fewest IRTs was selected for subsequent reinforcement. (In the case of ties, the bin corresponding to the shorter IRT was chosen.) The bird received reinforcement again only if it emitted an IRT that fell within the limits of this bin.

A second independent variable, in addition to λ , governed reinforcement frequency. This, the "probability of reinforcement" or $Pr(S_r)$, was included to reduce the number of reinforcements allowed by the IRT contingency just described. When $Pr(S_r) = 1$, all responses meeting the IRT contingency were reinforced; when $Pr(S_r) = 1/2$, about half of such responses

²This number varied from 64 to 152 during the early phases of the experiment. The number of IRTs saved was fixed at 150 during the seventh set of experimental conditions listed in Table 1.




Fig. 2. This indicates the way in which the IRT bins used to compute reinforcement are derived from an ideal exponential of given λ . The area under the curve is divided into 16 equal parts. The time interval spanned by each of these segments becomes an IRT bin. The rightmost bin has no upper bound. The ideal subject would fill these bins equally; the real subject puts fewer IRTs in some bins than in others. The subject is reinforced for producing an IRT that falls into the bin that contains the fewest IRTs. (Distribution starts at 0.8 sec, to exclude "double pecks." See text.)

were reinforced, and so on. This probability came into play each time a new IRT bin was selected for reinforcement. At that time, the computer used a random process to specify into which range must fall the IRT *preceding* the one potentially reinforced. With $\text{Pr}(\text{Sr}) = 1$, this preceding IRT could fall into any reinforcement bin. With $\text{Pr}(\text{Sr}) = \frac{1}{2}$, the preceding IRT had to fall into either (at random) one of the eight longer IRT bins, or one of the eight shorter IRT bins. With $\text{Pr}(\text{Sr}) = \frac{1}{4}$, the 16 reinforcement bins were split into four groups of four; the preceding IRT had to fall into one of these groups, selected at random.

This complex way of reducing reinforcement frequency was adopted to realize a potential extra push toward IRT variability. It favored random sequences of IRTs, where the basic reinforcement contingency only favored widely distributed values of IRT.

Four details of the reinforcement procedure remain. (a) The first response after reinforcement does not define an IRT in the usual sense, since eating time intervenes between responses. The IRTs for such responses were not determined and did not enter any IRT distribution involved in the experiment. Consequently, the first response after reinforcement never produced reinforcement. (b) When $\text{Pr}(\text{Sr})$ was reduced to $\frac{1}{4}$, long periods could elapse between reinforcements, so long, in fact, that birds might begin to extinguish. Hence, the following additional rule: if 20 min elapsed since a bird's last reinforcement, $\text{Pr}(\text{Sr})$ automatically became equal to 1; when a reinforcement occurred, $\text{Pr}(\text{Sr})$ reverted to its programmed value. In practice, this seldom happened. (c) As mentioned above, responses defining IRTs of less than 0.8 sec were never reinforced. It might have been possible, however, for such responses to be maintained superstitiously. If, for example, the bird frequently emitted double pecks, the first member of a pair could trigger reinforcement, and the second member might still be emitted before the bird began to feed. For this reason, if reinforcement was in progress when an IRT of less than 0.8 sec occurred, reinforcement stopped immediately. (d) At the start of a session, the IRT bin to be reinforced was determined from the IRT distribution carried over from the previous session. Thus, with regard to reinforcement contingencies, the birds worked, in effect, in one long, continuous session.

The data presented here were accumulated during the last $7\frac{1}{2}$ months of experimentation. During this period, λ varied from $\frac{1}{4}$ to 2, and probability of reinforcement $\text{Pr}(\text{Sr})$ varied between $\frac{1}{4}$ and 1. One or both parameters were changed when the birds had appeared to reach a stable level of responding under the preceding conditions. Table 1 shows the order in which these conditions were presented and the number of consecutive sessions under each condition. As the experiment

Table 1

Order of Experimental Conditions and the Number of Consecutive Sessions on Each (λp = programmed rate in responses per sec, $\text{Pr}(\text{Sr})$ = probability of reinforcement)

| λp | $\text{Pr}(\text{Sr})$ | Consecutive Days |
|---------------|------------------------|------------------|
| $\frac{1}{2}$ | 1 | 6 |
| $\frac{1}{2}$ | $\frac{1}{2}$ | 10 |
| 1 | $\frac{1}{2}$ | 18 |
| 1 | $\frac{1}{4}$ | 10 |
| 1 | 1 | 11 |
| $\frac{1}{4}$ | 1 | 8,* 11 |
| $\frac{1}{4}$ | $\frac{1}{4}$ | 13 |
| $\frac{1}{4}$ | $\frac{1}{2}$ | 18 |
| $\frac{1}{2}$ | $\frac{1}{4}$ | 15 |
| 2 | $\frac{1}{2}$ | 16 |
| $\frac{1}{2}$ | 1 | 15 |
| $\frac{1}{4}$ | $\frac{1}{4}$ | 19 |
| 1 | 1 | 12 |
| $\frac{1}{2}$ | $\frac{1}{4}$ | 13 |
| 1 | $\frac{1}{4}$ | 14 |
| $\frac{1}{4}$ | 1 | 14 |

*By accident, one session at $\lambda p = 1$ and $\text{Pr}(\text{Sr}) = 1$ was run on day 9.

progressed, it became clear that $\text{Pr}(\text{Sr})$ had relatively small effects. Hence, when replicating conditions, only the extreme values, $\text{Pr}(\text{Sr}) = 1$ and $\text{Pr}(\text{Sr}) = \frac{1}{4}$, were repeated. In considering the data, only these values are mentioned. The "first experimental series" refers to the first application of these values, paired with the various λ . The "second experimental series" refers to the replication of these conditions, which, it will be noted, came in a different order the second time around.

Definitions

Four variables, two of them already mentioned, are symbolized in the remainder of this paper. For convenience, they are defined here.

λp —Programmed *lambda*, an independent variable. It may be interpreted as the rate in responses per second of an ideal subject emitting responses randomly in time. However, as

may be seen from Fig. 2 and equation 2, λp refers to the rate with which responses occur, given that 0.8 sec has elapsed since the preceding response. In effect, time starts at 0.8 sec on the IRT continuum, and responses occurring within the 0.8-sec interval are not counted. (See also the definition of λ_0 below, and footnote 3.)

Pr(S^r)—Probability of reinforcement, an independent variable. Given that an IRT falls into the reinforcement bin with the fewest IRTs, the corresponding response will always be reinforced if $Pr(S^r) = 1$, it will go unreinforced about one-half of the time for $Pr(S^r) = 1/2$, and so on. For details, see above.

λ_0 = Obtained lambda, a dependent variable. A subject's rate, over an 80-min session, in responses per second, with the same omissions of time and responses described for λp . Since rapid multiple pecks count as single responses, λp reflects the same information as responses per second, with the contribution from double pecks and bursts sifted out.³

IRTs/OP—Interresponse-times per opportunity, a dependent variable. This is a transformation of a raw interresponse-time distribution. It estimates the conditional probability that a response will occur within a given time interval (say, from 3.0 to 3.5 sec) after a response. When applied to idealized random data, the transformation yields a horizontal straight line (Fig. 1), indicating that the probability of response is independent of IRT. The method of calculating IRTs/OP, as well as other peculiarities of the transformation, are discussed by Anger (1956).

RESULTS AND DISCUSSION

Response Rate

Cumulative records. Figure 3 shows cumulative records from the last day on each set of conditions in the second replication series. The curves are grouped in triplets, one curve from each bird. The first 2047 responses of the

session contribute to each curve, except where the total for the session was less than this. The cumulative records are not continuous; each successive point adds to the cumulative total the responses emitted in 20 sec. A dot appears next to the curve if one or more reinforcements were delivered during the associated 20-sec interval.

Figure 3 illustrates that the LF schedule generated a relatively stable rate within a session. The slope of the curves varies with λp , but not noticeably with $Pr(S^r)$. The cumulative records from different birds and different sessions were generally quite similar, although, as Fig. 3 suggests, variable performance sometimes accompanied the lowest rates of reinforcement.

Mean rate as a function of λp and $Pr(S^r)$. It has been shown that programmed lambda, λp , may be interpreted as the average rate of an ideal subject, while obtained lambda, λ_0 , represents the average rate of a real subject. Table 2 indicates the relation between these variables, together with the effect of $Pr(S^r)$. The data from the three birds are similar. They indicate that λ_0 varies substantially with λp , but little with $Pr(S^r)$. For $\lambda p = 1$, λ_0 is close to 1; for lower λp , λ_0 is somewhat higher than λp . In each case




Fig. 3. Cumulative response records from the last day on each λp and $Pr(S^r)$ combination in the second series. In each triplet of curves, the leftmost is from bird 007, the center from 593, the rightmost from 812. Curves in the top row all come from sessions with $Pr(S^r) = 1$; the bottom row from sessions with $Pr(S^r) = 1/4$. The curves are discontinuous, each point representing cumulation over an additional 20 sec. Reinforcement is indicated by a point next to the curve.

³An example will illustrate the difference between ordinary response rate and λ_0 . Suppose that four pecks come 0.4 sec apart, so that a total of 1.2 sec elapses from the first to the fourth. This group of pecks counts as a single response in the computation of λ_0 . In effect, the response lasts from the first peck until 0.8 sec after the fourth, for a total of $1.2 + 0.8 = 2$ sec. Since this 2-sec interval is not available for responses that count toward λ_0 , it is subtracted from the time used to compute λ_0 .

but one, $\text{Pr}(\text{Sr}) = 1$ is associated with a slightly higher λ_0 than is $\text{Pr}(\text{Sr}) = \frac{1}{4}$. (This difference is significant at the 5% level by a sign test).

Table 2

Adjusted Response Rate λ_0 In Responses Per Second As a Function of Programmed Rate λp and Probability of Reinforcement $\text{Pr}(\text{Sr})$
(Means over the last six days of each condition, including both experimental series)

| λp | Bird | | | |
|-------------|------------------------|------|------|------|
| | $\text{Pr}(\text{Sr})$ | 007 | 593 | 812 |
| .25 | $\frac{1}{4}$ | .35 | .30 | .36 |
| | 1 | .34 | .35 | .39 |
| .50 | $\frac{1}{4}$ | .53 | .56 | .59 |
| | 1 | .60 | .60 | .64 |
| 1.0 | $\frac{1}{4}$ | .89 | 1.01 | 1.02 |
| | 1 | 1.02 | 1.03 | 1.13 |

Rate across sessions. Figure 4 shows the observed rate, λ_0 , plotted for each bird over the last 84 sessions (the entire second experimental series). Pairs of numbers accompany each section of the graph; the top number is λp , the bottom $\text{Pr}(\text{Sr})$. Points cluster at about the same level for a given λp , regardless of $\text{Pr}(\text{Sr})$. However, there is somewhat more intra- and interbird variability at the lower $\text{Pr}(\text{Sr})$. With changing λp rate usually changed abruptly on the first day, followed by smaller adjustments over one or more subsequent sessions.

IRT distributions. Figure 5 summarizes IRT data by bird (rows), by $\text{Pr}(\text{Sr})$ (columns), by replication series (x and o symbols in each

graph), and by λp (three curves in each graph). These data are means from the last six sessions on a given set of conditions. The logarithmic ordinate makes it relatively easy to compare each curve with others and with the ideal exponential, which would be a straight line on these coordinates (Fig. 1).

All of the functions in Fig. 5 conform at least roughly to the straight line representing an exponential function. The consistent exception to this rule concerns responses falling in the first and second half-second bins. The level of these is quite variable and generally falls below the straight line described by the remainder of the data. (These short IRTs are discussed further below). The slope of the functions changes markedly with λp , but the different $\text{Pr}(\text{Sr})$ values (left and right cols.) have no clear effect.

Figure 5 indicates that considerable confidence can be placed in the form of the IRT data. Comparison of the rows shows that the three birds gave generally comparable results. A comparison of "x" with "o" curves in each plot shows that a given set of conditions gave about the same result each time it was used, despite the different order of conditions in the two series.

IRTs/OP beyond 1 sec. The comparability of the various sets of IRT data (Fig. 5) prompts the averaging of IRTs/OP functions to increase the visibility of any trends they may reveal. Figure 6 shows averages of these functions over all birds and both experimental series. Omitted from these functions are points for IRTs of less than 1 sec and points from the tails of the distributions where the average included quotients with zero denominators.

As in Fig. 5, the most striking feature of the IRTs/OP curves is the separation of functions for the different λp values. The functions associated with $\text{Pr}(\text{Sr}) = 1$ are roughly constant, except for higher values in the 1-2-sec IRT range. With increasing IRT, IRTs/OP associated with $\text{Pr}(\text{Sr}) = \frac{1}{4}$ tend to fall below those for the higher reinforcement probability. The raw data reveal that this difference was due largely to the fact that many more very long IRTs (greater than 10.5 sec) occurred at the lower $\text{Pr}(\text{Sr})$. This is consistent with the slightly lower λ_0 values associated with $\text{Pr}(\text{Sr}) = \frac{1}{4}$ (Table 2). It may be tentatively concluded that fewer reinforcements increase the number of very long IRTs some-




Fig. 4. Response rate, given by obtained lambda (λ_0) over the last 84 experimental sessions. Each segment of the curve is labeled with the values of the independent variables associated with that segment. In each pair, the programmed rate λp is on top, and the probability of reinforcement $\text{Pr}(\text{Sr})$ below.

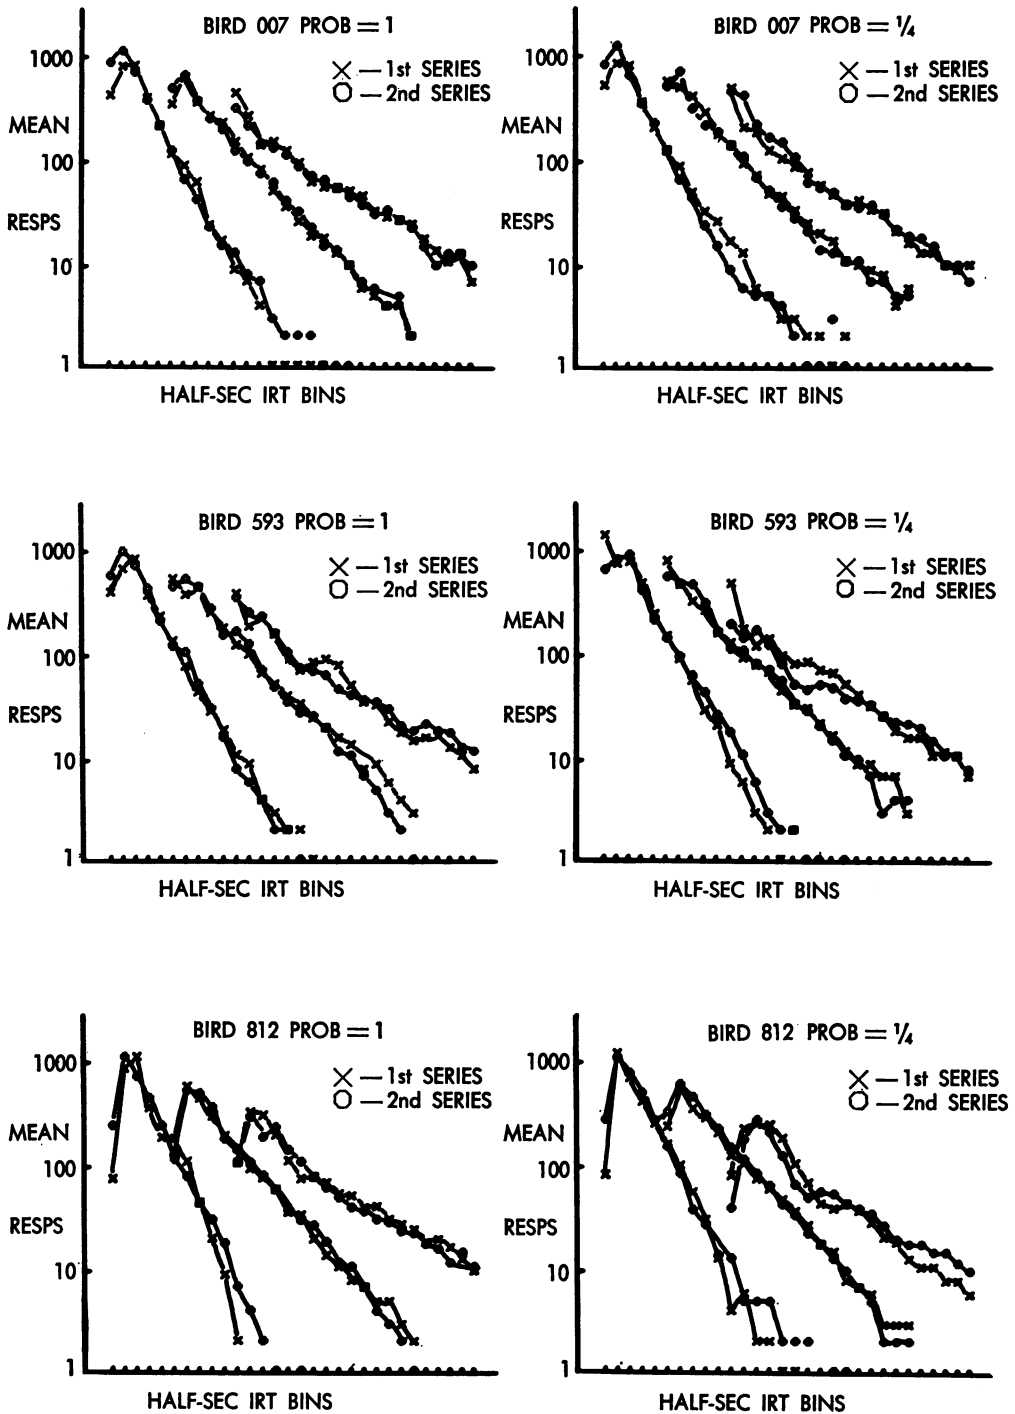


Fig. 5. Mean IRTs over the last six sessions under each experimental condition. Means appear for 12 sets of sessions, two sets for each of the following combinations of independent variables: $\Pr(S^*) = 1$, with $\lambda p = 1$, $1/2$, and $1/4$; $\Pr(S^*) = 1/4$, with $\lambda p = 1$, $1/2$, and $1/4$. The three curves in each graph correspond to the three λp values, from left to right: 1, $1/2$, $1/4$. The curve for $\lambda p = 1/2$ has been displaced to the right five units; for $\lambda p = 1/4$, 10 units. To avoid confusing overlap, the catch-all IRT bin for IRT greater than 10.5 sec has been omitted from the graphs.

what, but do not markedly affect the IRT distribution when the bird is active.

An apparently minor feature of the procedure may have contributed to the relatively high IRTs/OP values at short IRT. As described earlier, reinforcement was set up for responses falling into the low IRT bin. A sampling of the data indicated, however, that from 10% to 25% of the time, two or more bins were tied for low. In such cases, the computer always chose the shorter IRT bin for reinforcement. This set a bias in favor of short IRTs that may have been reflected in the IRTs/OP curves.

IRTs/OP at short IRT. Blough (1963) suggests that the pigeon's short IRTs have a special character that justifies their separate treatment. The present study supports this view. Figure 7 provides the basis for this statement. In it appear plots of IRTs/OP by 0.1-sec bins for the three birds during the first 2 sec after the preceding response. These curves differ in two ways from the data so far discussed: they show large individual differences among subjects at IRTs less than about 0.7 sec, and they indicate that λ_p has little or no effect on response probability in this region. IRT data across series (*cf.* Fig. 5) indicate that response probability in this short IRT region shows the largest session-to-session variability. Such IRTs showed no sign of diminishing, though responses terminating IRTs of less than 0.8 sec were never reinforced.




Fig. 6. Mean IRTs/OP over the last six days of each condition, of both experimental series, computed for half-second IRT bins. Points below 1 sec are missing (*cf.* Fig. 7); long IRT points are missing where their meaning is questionable because few responses contributed to them.

It appears that these short IRT responses are patterned in a stereotyped way peculiar to the subject. They are perhaps best considered as part of the pigeon's response topography, the double peck being simply a variant of the peck. They would seem worthy of little attention, were it not for the fact that they constitute a very large proportion of some birds' recorded output, and appear to change their proportion as response topography drifts. Thus, these annoying extra pecks may considerably increase variability across subjects and from time to time.

Figure 7 suggests further that this response stereotypy has effects extending to IRTs as great as 2 sec. For example, a damped oscillation with a period of about 0.4 sec is evident in the data of bird 593, with progressively less marked IRTs/OP peaks at 0.4, 0.8, 1.2, and 1.6 sec. As Blough (1963) noted, such periodicity is probably due to a basic IRT of 0.4 sec, the longer peaks resulting from abortive responses that fail to operate the response key. Whatever their cause, such periodicity interferes with the operation of the schedule, which does not take into account this factor in allocating reinforcements within the short IRT period. It was primarily this situation that prompted adoption of the 0.8-sec dead time after each response, within which no reinforcements were given.

If very short IRTs are indeed the result of special topography, they may be controllable by careful key training. Even without them, the short-IRT region presents a problem for the schedule. Real subjects can never approach the ideal emitter here, since they require a finite minimum IRT. This minimum surely varies, so the assumption of the present schedule, that the bird becomes active at exactly 0.8 sec, is at best a gross simplification. The situation might be handled by adjusting reinforcement bin sizes in the short IRT region to take account of rising response probability there. The data of bird 812, a bird with clean topography and few double pecks, may suggest the sort of distribution to be taken into account (Fig. 7).

Reinforcement Distributions

Anger (1956) suggested that the conditional probability of response in a given IRT bin (estimated by IRTs/OP) is proportional to the relative number of reinforcements per unit

time delivered in that bin. On this hypothesis, a constant conditional probability across IRT bins should be accompanied by a constant number of reinforcements per bin. The LF




Fig. 7. Mean IRTs/OP over the last six days of each condition of both experimental series, computed for 0.1-sec bins. Data over the first 2 sec are shown. For IRTs less than 0.7 sec, note large individual differences and the lack of effect of λp .




Fig. 8. Distributions of reinforcements by IRT, in half-second bins, over the second experimental series. Percentages are used to enable comparison between reinforcement probabilities; were absolute values plotted, the curves for reinforcement probability = 1 would be much higher than for reinforcement probability = $1/4$. Each point represents a mean over the last six days of a given experimental condition.

schedule was intended to yield a constant conditional probability, and it succeeded to the extent shown in Fig. 6. Therefore, reinforcements per bin might be expected to be approximately constant.

Figure 8 shows distribution of reinforcements, combined across birds, for the last six days of the various conditions in the second experimental series. (Bear in mind that the program prohibited reinforcements for IRTs less than 0.8 sec). The curves for $\lambda p = 1$ depart most strikingly from the horizontal straight line that would indicate constancy across bins. The steeply falling curve here is more compatible with the rule that the number of responses in a given bin, not the conditional probability of response, is proportional to the number of reinforcements in that bin. This rule was rejected by Anger, and further evidence against it has since appeared (Catania, 1962; Herrnstein, 1964). For $\lambda = \frac{1}{2}$ and $\frac{1}{4}$, the reinforcement distributions are not proportional to responses, yet they do not achieve the constancy required by Anger's rule.

The variability in the reinforcement distribution data is large, but the data suggest that no simple rule, such as "equal reinforcement per IRT" or "equal reinforcement per response" applies under all conditions. They suggest rather, that the pigeon's rate is not governed entirely by reinforced IRT, powerful as this factor appears to be. To force a bird to high rates (many short IRTs), it may be necessary to reinforce short IRTs heavily; to force a bird to slow down, long IRTs must be heavily weighted.

Interresponse Dependence

Even had the present subjects generated purely exponential IRT distributions, it could not be argued that their responses were independently emitted. The latter implies the former, but not the reverse. Additional tests of sequential ordering of IRTs are necessary. The last 150 IRTs (omitting those shorter than 0.8 sec) were retained from each experimental session in the second experimental series, and two simple tests of serial dependence were run on these samples. The first was a runs test (Dixon and Massey, 1952). IRTs were dichotomized into short and long and the number of runs in the resulting sequence was computed. This statistic was com-

pared with the number of runs to be expected on the basis of IRT independence. Of 108 samples tested, 24 had too few runs (significant at the 5% level or better) while three had too many runs. Eighteen of the series with too few runs were from one bird, 812.

In this analysis, too few runs indicates excessive repetition of IRTs, long following long, short following short. Unfortunately, it reveals no more about the patterning of the runs. In scanning the data, no particular patterning was detected.

The second test was a chi-square analysis of data from the last six days on each experimental condition. For this analysis, each pair of IRTs in a sequence was categorized into a two by two table: short-short, short-long, long-short, long-long. Of chi-squares from 30 such tables, 12 proved significant at the 5% level or better. This result confirmed the runs tests, indicating some degree of serial dependence expressed as repeated long or short IRTs.

Though interresponse dependence did not appear to play a major role in this study, its presence constitutes response control that may be undesirable. More might be done to control it by basing reinforcement upon IRT sequence in addition to current IRT, the tactic used to a minor extent in the present study, when $Pr(Sr)$ was reduced (see Procedure).

Uses and Implications of the LF Schedule

The LF schedule largely achieved its main purpose: control of mean rate while diversifying local behavior. Changes in the rate constant λp altered response rate rapidly, precisely, and reliably enough that an experimenter may use the schedule to generate approximately the rate desired, at least within the range tested here. Thus, the schedule appears useful where a predictable non-stereotyped rate baseline is required, given that this baseline is not itself supposed to be sensitive to independent variables. Two examples may clarify such applications. (1) The LF schedule is currently being used in a steady-state discrimination experiment. It maintains stable responding to a discriminative stimulus, while responding to other stimuli varies. (2) The schedule might maintain behavior to one key in a two-key situation, where responding to the second key was of interest. For example, one might test the implications of quantitative formulations involving relative response and

reinforcement rates (e.g., Catania, 1962) since the schedule permits independent control of these two variables.

It would not be appropriate to use the schedule in its present form to study the effects of variables (e.g., punishment, deprivation) upon responses that were directly under schedule control. The schedule opposes rate changes, like the governor of an engine; it successfully opposed the effect of a large change in the probability of reinforcement (assuming that such a change might otherwise have been effective). Probably variants of the LF schedule could be devised that would make it useful for such purposes. For example, the programmed rate constant λp might be put under the subject's control, letting it drift upward if pressure toward a higher rate developed (too many short IRTs) and downward if the subject tended to slow down. As in other systems involving feedback, such a procedure could lead to a stable equilibrium, oscillation, or a flight to one extreme (extinction) or the other (maximum rate). The outcome would no doubt depend upon the details of the feedback procedure.

The present study adds weight to the already large body of evidence (c.f. Ferster and Skinner, 1957) that behavior at the time of reinforcement powerfully determines the performance generated by a schedule. The results do not imply, however, that reinforced interresponse time is the only effective variable controlling the performance. The effectiveness of this variable was abnormally enhanced in the present study because it was specifically used to oppose the effects of other variables. Yet, while granting that other variables may control response emission, one should always be careful to examine the possibility (as did

Anger, 1956) that a variable gains part or all of its effect indirectly, by governing the distribution of reinforced interresponse times.

REFERENCES

- Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *J. exp. Psychol.*, 1956, 52, 145-161.
- Blough, D. S. Interresponse time as a function of continuous variables: a new method and some data. *J. exp. Anal. Behav.*, 1963, 6, 237-246.
- Blough, D. S. and Millward, R. B. Learning: operant conditioning and verbal learning. *Ann. Rev. Psychol.*, 1965, 16, 63-94.
- Catania, A. C. Independence of concurrent responding maintained by interval schedules of reinforcement. *J. exp. Anal. Behav.*, 1962, 5, 174-184.
- Clark, W. A. and Molnar, C. E. The LINC: A description of the Laboratory Instrument Computer. *Ann. N. Y. Acad. Sci.*, 1964, 115, art. 2, 653-668.
- Dixon, W. J. and Massey, F. J. *Introduction to statistical analysis*. New York: McGraw-Hill, p. 254, ff., 1951.
- Ferster, C. B. and Skinner, B. F. *Schedules of reinforcement*. New York: Appleton-Century-Crofts, 1957.
- Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *J. exp. Anal. Behav.*, 1964, 7, 27-36.
- McGill, W. J. Stochastic latency mechanisms. In *Handbook of mathematical psychology*. Vol. I. Luce, R. D., Bush, R. R. and McGill, W. J. (Eds.) New York: Wiley, 1963.
- Mueller, C. G. Theoretical relationships among some measures of conditioning. *Proc. Nat. Acad. Sci.*, 1950, 36, 123-130.
- Sidman, M. *Tactics of scientific research. Evaluating experimental data in psychology*. New York: Basic Bks., 1960.
- Skinner, B. F. Superstition in the pigeon. *J. exp. Psychol.*, 1948, 38, 168-172.
- Weiss, B. and Laties, V. G. Reinforcement schedule generated by an on-line digital computer. *Science*, 1965, 148, 658-661.

Received December 28, 1965